Don't Patch the Researcher, Patch the Game: A Systematic Approach for Responsible Research via Federated Ethics Boards

Alexandra Dirksen
Technische Universität Braunschweig
Braunschweig, Germany
a.dirksen@tu-braunschweig.de

Sebastian Giessler
University of Tübingen
Tübingen, Germany
sebastian.giessler@izew.uni-tuebingen.de

Hendrik Erz
Linköping University
Norrköping, Sweden
hendrik.erz@liu.se

Martin Johns
Technische Universität Braunschweig
Braunschweig, Germany
m.johns@tu-braunschweig.de

Tobias Fiebig
Max-Planck Institut für Informatik
Saarbrücken, Germany
tfiebig@mpi-inf.mpg.de

Abstract

Advances in technology enable researchers to incorporate increasing amounts of user data into their research studies. However, direct and indirect involvement of users and/or their data may negatively affect users, e.g., via privacy violations or by exposing them to security vulnerabilities. Therefore, researchers must follow research ethics principles and address their work's ethical and social implications. This process is often guided and monitored by institutionalized bodies and ethics committees, usually embedded in research institutions and publication venues. Recent events in Computer Science (CS) research on the boundary between human/social and technical research highlighted limitations in this process. Ethical reviews are effective when human involvement is evident. However, they are challenging for projects where human involvement is indirect or not the primary focus, as is mostly the case in CS. This led to cases where work was, e.g., initially authorized by the local ethical review yet later found to be unethical by the community, leading to the retraction of those papers. However, the potential harm remains done. In this paper, we take a systematic approach towards ethical reviewing procedures for CS. We do not advocate patching the current system, enforcing it legally, or creating a uniform norm framework. Instead, we take a step back and revisit the system as a whole: We first investigate the current practices of ethical reviewing, which heavily rely on individual responsibility and ex-post (after-the-fact) reviews and retractions. From that, we propose a novel FEB Federation to address systemic shortcomings of the ethical reviewing procedure in CS.

CCS Concepts

• **Security and privacy → Human and societal aspects of security and privacy**; • **Social and professional topics → Codes of ethics**.

NSPW'24, June 16–19, 2024, Bedford, PA USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/XXXXXXX.XXXXXXX

Keywords

research ethics, security & privacy research, institutional review boards, fairness

ACM Reference Format:
Alexandra Dirksen, Sebastian Giessler, Hendrik Erz, Martin Johns, and Tobias Fiebig. 2024. Don't Patch the Researcher, Patch the Game: A Systematic Approach for Responsible Research via Federated Ethics Boards. In *Proceedings of The New Security Paradigms Workshop (NSPW'24).* ACM, New York, NY, USA, 16 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Computer science-related research has long since departed from a purely technological perspective. Instead, computer scientists, starting with the human-computer-interaction (HCI) research community, but also beyond, started to consider 'human and societal factors' in their work [8]. How the human factor materializes in studies differs based on the scope of the specific work. This ranges from directly related work, e.g., usability focused studies [26, 45], or studies in which social effects are studied [13, 78, 82], to work like censorship measurements [28, 36, 69] or mathematics [29].

Naturally, studies in some way concerned with human or societal factors can raise moral and ethical concerns. This may be due to the results of research work and its potential applications (e.g., systems that directly influence consumers, like calculating credit scores or compiling news feeds in social media), or by harming participants and others while conducting the study.

However, morals can be perceived as a subjective issue that can be influenced by a society's values, such as national and cultural perspectives and values and the standards established in specific academic disciplines. What is considered moral in one part of the world may not be regarded as moral in another, and acceptable research in one community may be considered moral in another, and vice versa. Consequently, ethical governance varies widely in the CS community to ensure that CS research dealing with such data complies with some ethical standards. Companies, governments, and organizations who claim leadership in the discussion of technology ethics, like the Association for Computing Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE), focus on non-binding voluntary guidelines without any enforcement mechanisms in case of scientific misconduct [3].

However, in practice, these structures are not always effective and sometimes even miss aspects of a study for which it later receives critical commentary, either from within the community, the wider public, or the communities from which the participants were recruited. This is especially frequent for studies that involve users in experiments, e.g., by indirectly measuring users' behavior during daily internet use [59, 77], or by directly having them react to certain events. Please note that to promote blame-free research, we will refrain from passing judgment on the effects and motivations of these studies and will refrain from explicitly referencing them when possible. Instead, we refer to related work discussing such cases in detail, e.g., [33, 41, 55, 75, 84] and will only describe one example scenario in detail to illustrate these mechanics.

Universities carry different denominations like Research Ethics Boards (REB), Research Ethics Committee (REC), Ethical Review Boards (ERB), or the U.S.-based Institutional Review Boards (IRB), which offer ethical oversight and/or governance. For brevity, we refer to them all as **ERB**. However, for clarity and to distance our proposal in this paper from the current systems in place, we later (Sec. 5) introduce to the acronym **FEB** for such boards.

Example Case: In April 2021, a paper accepted to the IEEE Symposium on Security & Privacy received widespread attention from the CS community [89]. The paper aimed to investigate whether malicious actors could submit patches introducing vulnerabilities to the Linux Kernel project. To this end, the researchers submitted patches for known issues to the Linux Kernel Mailing List (LKML), which additionally introduced a new vulnerability. If a maintainer accepted such a patch, the authors would notify the maintainers about the hidden vulnerability and the study, therefore "un-blinding" the specific participants.

However, due to the deception design, this study lacked informed consent. Therefore, it was received with discontent by kernel maintainers, see, e.g., a public email exchange with the researchers [50]. Following increased public attention to this mail exchange, the researchers' institution retracted a press statement welcoming the paper's acceptance to the IEEE Symposium on Security & Privacy [23]. Furthermore, the paper was also retracted from the proceedings due to a lack of following ethical principles, namely the well-established principles of informed consent, non-maleficence, and non-deception [43].

Still, according to the authors of the paper, the ERB of their research institution had reviewed the project and *"determined that this is not human research"* [89], and therefore granted a formal waiver of the exemption. However, the experiment's outcome required the involvement of humans—a Linux Kernel maintainer who had to reply to an email. The paper's underlying research question was clearly, "Do humans spot malicious code?" The unit of analysis (UoA) was individual people. Nevertheless, this connection may not have been apparent to an ERB that is usually concerned with more *direct* human subject involvement, e.g., when visual stimuli are applied to participants, leading to the observed outcome.

In response to this and the general issue of making ethical judgments one way or another as a reviewer, the security and privacy community started to implement 'Ex-Post' ethics committees in several major venues. These committees are staffed by community members and can be called in by reviewers if a submitted work raises ethical questions. The ethical subcommittee then assesses the case, discusses it with the authors, and may even solicit additional out-of-band information from the authors, e.g., on submitted material to an ERB. This process should bring more consistency into the ethical review process, i.e., prevent papers from being rejected based on the subjective ethics of individual reviewers while also being able to identify more broadly ethically challenging work in the form of community self-regulation. See, for example, the IEEE Symposium on Security & Privacy '23 CFP [73]. Naturally, however, this approach still suffers from it's 'Ex-Post' nature, and - more generally - from the issue of hindsight bias.

Hindsight: There is a notable lack of clarity about what requires an ethical investigation and/or the prior consent of participants. An ERB, only following state-of-the-art guidelines for ethical research, e.g., [5], would possibly authorize a measurement project on Internet censorship if it does not require storing any user data. However, there is a chance they do not consult other expert sources due to a lack of domain knowledge. In that case, they tend to neglect the risks when those measurements are carried out via resident proxies in countries with politically motivated censorship.

In retrospect, especially after participants have complained [30], judging whether something may not have been ethical is often easier. Sometimes, such concerns are also raised during conference reviewing, when more people with different perspectives assess conducted research. However, since the experiments have already been carried out and cannot be undone at the time of the submission, those responsible for the publication have no choice but to either withdraw the paper [90] or publish it with an appropriate statement [77]. Still, at that point, harm has already been done. Instead, these critical issues should be identified and addressed during the initial design and review process. Even in hindsight, it is sometimes impossible to trace the source of the misjudgment or whether it was avoidable. Especially if one takes a just culture [21] perspective. Hence, such cases demonstrate that the present structures of ethical reviewing are insufficient to create good outcomes. Good outcomes, here, are understood as outcomes where minimal harm is being done, and the research itself holds up to be considered ethical by relevant stakeholders.

Challenges in Ethical Review: Given the potentially varying backgrounds and ethical frameworks researchers operate under, we must address this problem systemically, i.e., not focused on individual institutions' and researchers' responsibilities:

- The general system of handling ethics in CS must ensure that ethical reviews are done *before* the experiments are conducted.
- The system must ensure that researchers are necessitated to provide the required information for the ethical review.
- Despite varying ethical perspectives, the ERB needs to have the necessary domain knowledge to estimate where ethical problems can occur.
- While the system may rely on all participants acting in good faith and being unbiased, it should be resilient against human error and mistakes.

We argue that our current ethical review approach is structurally inadequate for addressing the ethical challenges of CS research. This inadequacy stems from a lack of responsibility or domain knowledge and a belated involvement of oversight bodies, limiting their impact on the research projects. As the direct or indirect involvement of social data in CS research becomes more widespread [11] and stakeholders widely recognize ethical issues [4, 63], uniform and robust oversight procedures are needed. Such a procedure for ethics review can lead to better oversight and accountability in CS research, more transparency about how user data is used and interpreted, and structures where people affected by research can act accordingly. This serves the protection of human subjects and ensures ethically responsible research.

Scope: Towards this goal, we rethink the CS community's approach to ethical reviewing in CS and introduce a new, unified system. However, we do not propose a normative framework for what is or is not ethical in CS research, nor do we advocate a legal basis for the review. Instead, we strive to implement a framework that ensures good outcomes (harm reduction and rough consensus) and propose a procedure and interaction protocol for enabling ethical reviewing. To further illustrate our framework, we focus on the ever-growing field of Security and Privacy research as part of CS, a discipline concerned with protecting users' privacy and property in the digital space. Still, our framework generalizes beyond this illustrative scope.

Contributions: In summary, we make the following contributions:

- We systematically analyze the current limitations of ethical review of research on the boundary between human and social factors with CS.
- We propose a unified procedure for the ethical review of research projects that does justice to a heterogeneous landscape of ethical values in different cultural and scientific communities by adopting a system-oriented rather than a normative approach to ethics.
- We analyze the participation incentives for universities and research institutions and outline our approach's implementation challenges that must be overcome to implement it in practice.

Structure: The remainder of this paper is structured as follows. First, we will provide background, terminology, and broader context on ethical review and standards in Section 2. Next, in Section 3, we conduct a study of the mandated ethical review procedures in several countries across the globe. Based on our observations, we derive requirements for our framework in Section 4, based on which we derive the actual framework in Section 5. We then discuss practical and implementation considerations in Section 6. Finally, we conclude in Section 7.

2 Background

In this section, we first introduce terminology used throughout the paper, including the concept of "harm". We then provide a perspective on how ethics is being handled in CS research, discussing the current de-facto standards in the form of the Belmont report [37] and its updated successor, the Menlo report [5]. This also includes a description of how current ERBs relate to research ethics and the role they should fulfill in the social infrastructure of science.

2.1 Terminology

Within our work, we use various terms commonly used in ethical discussions. To make these terms accessible and defined, we briefly describe them here.

2.1.1 Ex-ante and ex-post. When discussing the ethics assessment for a given research project, it is essential to have clear terms for when the ethical review takes place. Naturally, two options exist: either before or after the first experiment or measurement has been conducted. We refer to these two as *ex-ante* and *ex-post*.

Ex ante. An ex-ante ethics review is conducted before a research project conducts any work concerning the evaluated research, e.g., collecting data, performing experiments, or otherwise interacting with participants. Ex ante reviews are commonly conducted by ERBs in academic institutions by mandated legislative oversight and regulatory bodies. Similarly, communities that have adopted "study pre-registration" [67] sometimes integrate ethical review (or provide proof this has been done) along with the pre-registration process. An ex-ante review may have a preventive effect, i.e., it occurs at a point in time when ethically unfeasible research can still be prevented.

Ex post. An ex-post ethics review occurs when the research has already been completed. Conferences or journals can integrate it into the peer-review process and make it explicit. It can also occur implicitly, e.g., based on uproar in civil society or the investigated population after the published study. It may then be explicitly relegated to, e.g., the scientific society/community review board in which the concerned work appeared. In any case, despite potentially punitive or restorative actions, ex-post reviews do not have any leverage to prevent harm.

2.1.2 Ethics / Ethical. Philosophical ethics is concerned with the distinction between right and wrong actions, what should or should not be done for what reasons[1]. Research ethics investigates the ethical dimensions of scientific research and research practices [71, 83]. This includes traditional topics of Responsible Conduct of Research (RCR) like research misconduct, authorship, conflicts of interest, data management, animal welfare, and human subject protections, but goes beyond RCR in analyzing the external consequences of the application of scientific research, e.g., the impact on society [83]. Further, research ethics is commonly concerned with whether a specific research project should (not) be done, whether a project is safe, what constitutes a risk, etc. Thus, when we discuss ethics in this work, we refer to this specific subset of ethics related to research practices or, even more specifically, along our example ethics of Computer Security [56].

However, it is crucial to be aware that - unlike the physical and social sciences - research ethics does not provide descriptive knowledge [71]. Thus, the field of research ethics is an advisory authority – it cannot establish values and norms on behalf of scientific research and researchers. Since CS has no legally binding structures for conducting human-subject research, ethical oversight is based on two pillars: On the individual mindfulness of the researcher to adhere to ethical principles and guidelines of their profession (see, e.g., IEEE Guidelines), and institutionalized oversight bodies,

[1] See also "Applied Ethics", e.g., at https://iep.utm.edu/applied-ethics/.

the ERBs. The objective is to prevent potential harm to research subjects. We argue that both pillars of ethical oversight to ensure RCR in CS systematically fails to meet this objective.

2.1.3 Rough Consensus. Ethical considerations can be rooted in a variety of philosophical traditions. Examples from Western philosophical thought include deontology, utilitarianism, and virtue ethics, but frameworks based on African philosophy (Ubuntu ethics) and Confucian ethics are also well established. Principle ethics are also frequently used in research, prominently argued by Beauchamp and Childress in Principles of Biomedical Ethics with their concept of globally converging moral norms [7]. Since the specific methodology for judging decisions and ethical frameworks is contested and the subject of fierce discussion, we want to refrain from prescribing any moral frameworks. We are discussing the governance and oversight mechanisms employed to ensure ethical research, not how to judge what research should be considered ethical. Each ethical-normative system has its focus and employs fundamentally different approaches. None of the systems is worse or better than the others, but a "one size fits all" ethical approach is neither possible nor desirable. For an overview of how different ethical theories view technological problems, we refer to [62].

Rough consensus is a concept popularized by the Internet Engineering Task Force (IETF), a body that standardizes Internet protocols. Given the Internet's environment, which often involves diverging needs and opinions, e.g., among different vendors favoring different implementations, the IETF adopted the concept of 'rough consensus' to establish whether a document or proposal is adopted. Rough consensus, as detailed in RFC7282 [70], describes a procedure focused on *"lack of disagreement"* instead of *"presence of agreement"*, aiming for *"[…] all issues [being] addressed, but not necessarily accommodated"*. Therefore, it is a mechanic aimed at forming progressing and adaptive standards within a community with diverging viewpoints on individual aspects of these standards.

2.1.4 Harm. Defining harm is generally difficult[2]. One approach is to consider objective harm, which occurs when an action has a universally recognized negative impact on an individual. This impact could be physical, psychological, social, or moral, often disrupting the individual's well-being or ability to thrive. In the context of research ethics, harms can vary widely. They might include failures to secure informed consent, breaches of confidentiality concerning personal details like sexuality or political beliefs, financial losses, psychological trauma, or even physical injury, including death.

However, harm deeply interacts with what the harmed individual perceives to be harmful [16], i.e., it is a subjective property. Hence, when discussing harm, we naturally include all forms of subjective harm as the organic superset of objective harm. This increases the difficulty of preventing harm, as researchers must anticipate what causes harm. However, it also clearly characterizes responsibilities during the ethical review, i.e., that it has to entail reflecting on possible subjective harm. It must include appropriate domain knowledge to conduct this reflection. Depending on the population on which research is conducted, this may directly necessitate the involvement of community members. This is especially true as research involves marginalized communities or communities under risk [8].

[2]In recent years, the scope of research ethics has been broadened to include factors such as the environment, animal welfare, and social welfare. As we are talking here specifically about research ethics in the sense of human subjects research in (CS), we use the term to avoid personal harm.

2.2 On the Ethical Aspects of Research

Moral philosophy and ethics, or philosophical ethics in general, is the scientific endeavor of considering, evaluating, and discussing moral questions [49]. In short, it's about deciding what is right and wrong and how to differentiate these positions. Narrowing this question to research ethics means answering the question of which research should or should not be conducted. Thus, research ethics can only be an advisory authority that cannot establish values and norms on behalf of science and researchers [72]. The biomedical and behavioral sciences primarily use elaborate ethical frameworks to navigate those challenges [52]. In the case of CS research, Kohno et al. (2023) worked out an instructive paper on the normative foundation of ethical and moral discussions in computer science and security research. We want to pick up on that work and propose an actionable suggestion to implement the issues Kohno et al. outlined. An important takeaway from this paper is that there is no unified position on deciding on ethical challenges in computer science. Approaches informed by ethical reasoning may result in different conclusions by different people on similar or comparable ethical challenges [49].

However, we want to avoid justifying the adherence to research ethics anew in this paper. Therefore, we assume two positions: First, research ethics is an integral part of science, and it is necessary to adhere to the ethical conduct of research. Second, according to van Heerden et al., we assume that in CS research, the fundamental principles of conducting ethical, social research remain the same [74], i.e., CS research should be subject to the same scrutiny and oversight as other scientific disciplines. Extended work on research ethics in CS has already been made by Macnish, Buchanan, Christen, and others [11, 15, 55], supporting this thesis. However, the ethicization of technological conflicts implies shortcomings and increased risk awareness concerning large-scale research using personal data in CS and neighboring disciplines [46]. Consequently, institutions such as the ERBs are an essential part of the social infrastructure of science, whose task is to oversee adherence to RCR and research ethics.

Nevertheless, as already described, there have been examples of recent CS research projects where ethical issues either have been neglected or research has been conducted that might not have received ERB approval, or the decisions of ERBs were questioned afterward by the conference reviewers [17, 47]. It should be noted that this research was by no means illegal. This paper focuses on ERBs as part of institutionalized scientific governance structures directly influencing the process of knowledge production. Thus, cases that rely heavily on obviously illegally obtained personal data, like the Internet Census 2012 outlined in the Menlo Report [27], or non-binding Ethics Committees or guidelines, are excluded.

2.3 On Guidelines and Ethical Principles

Moral philosophy and ethics are scientific fields that have existed for several thousand years. Ethical frameworks offer valuable perspectives on technology and justice. The three classical approaches in Western philosophy, virtue ethics, deontology, and utilitarianism, offer various criteria for moral behavior and justice. They provide

guidance on what ought to be done or omitted to promote human welfare, minimize harm, and create justice. As said in Sec 2.2, these decisions may vary over geographical or temporal divides. Regarding research ethics, common principles have emerged that are mainly followed in human subjects research. As argued by Buchanan et al., [11], research ethics have two major obstacles to overcome in CS research:

- The established standards for human subjects research have evolved from medical, biomedical, and behavioral research perspectives, particularly within U.S. regulatory frameworks. These regulations were crafted with these fields in mind without explicitly considering technology-focused research. Consequently, there exists a significant challenge in aligning CS research with these established standards for human subjects research [10, Ch. 74]).
- Computer Science is partially a formal science focusing on mathematical optimization, algorithm development, logic, and computer architecture, among other areas. Only a subset of the discipline can be considered empirical research that involves human subjects. The potential for harm in this context is often more abstract than in medical, biomedical, and behavioral sciences. For instance, there is an intuitive distinction between a medical intervention on a person and the exposure of a deanonymized IP address. This distinction leads to two phenomena: computer science researchers often lack structured training in handling human subject research, and existing ethical governance institutions may not possess the necessary domain knowledge to assess potential harm to individuals adequately.

Common rules and self-governance are limited in effectiveness, particularly concerning emerging technologies. Specific technologies cannot be tightly regulated within a research context because the landscape constantly evolves, making it a "moving target". While numerous guidelines and recommendations exist for RCR [3], the absence of strict enforcement mechanisms still leaves room for potential harm.

In CS research, ethical principles often refer to the Menlo Report and its precursor, the Belmont Report [5, 37]. They formulate standards of conduct for human subjects research in the US, and the Menlo report, in particular, was formulated with CS research and development in mind. The following ethics principles were formulated in the Menlo Report and provide an anchor for ERBs in the US, whose task is to balance these principles to protect human subjects against legitimate research interests.

Respect for the autonomy, privacy and dignity of individuals and communities: This principle refers to the direct handling of individuals and individual data. This stretches from privacy, confidentiality, and copyright issues to informed consent, a standard for human subjects research [74]. In the context of online data collection, this refers to the problem of whether subjects are informed by the researcher about data collection, monitoring, and what data is being collected [87].

Beneficence: Beneficence refers to weighing relative risks for the subjects against the research benefits. This principle is often translated into the rule "do not harm", which requires maximizing probable benefits and minimizing possible harms. In the context of online data collection, this means, for example, preserving anonymity and protecting identity. These harm and benefit risks must be systematically assessed [87].

Justice: "Each person deserves equal treatment, and research benefits should be fairly distributed according to individual need, effort, societal contribution, and merit. Selection of subjects should be fair, and burdens should be allocated equitably across impacted subjects". One challenge for this principle is sampling bias due to self-selection, i.e., that the demographics of global Internet use do not match demographics in terms of sociodemographic factors such as gender, age, and socioeconomic status [5, 87].

Respect for law and public interest: Engage in legal due diligence; Be transparent in methods and results; "Be accountable for actions" [5].

However, due to the characteristics of CS research, it is becoming increasingly challenging to implement these principles. ERBs are often confronted with new ethical questions when overseeing CS research. What ethical obligations do researchers have to protect subjects' privacy in public Internet spaces? How is informed consent obtained and recorded by research subjects when conducting large-scale social media research? Is deception an acceptable online norm or harm [11]? In our conclusion, we argue that ERBs are ill-equipped to tackle the challenges of CS research. The lack of a unique *tradition* of Research Ethics in CS is further aggravating this issue [55].

3 ERB Procedures Around the Globe

As mentioned in Section 2, what *is* considered the correct moral choice differs by location, community, and society. Hence, ERB structures and procedures also vary across countries. Scientific research in CS is a global endeavor. Conferences and journals receive submissions from starkly different political, legal, and cultural environments and generally don't discriminate based on local legal systems. Ensuring at least some sort of minimal consensus on conducting responsible research is crucial. In this work, we take a qualitative perspective on ERB setups, i.e., we provide diverse examples of ERB implementations around the world but do not claim exhaustiveness. For an exhaustive perspective, we refer to existing studies that deal empirically with ERBs [2, 48, 60, 86].

Therefore, this section focuses on representative ethical governance structures that vary in their institutional and legal requirements: the United States, the European Union, China, India, and, as a sample, Germany and Sweden, among others. We will also briefly introduce the ethical governance structures of India and China and touch on independent/commercial ERBs. We have selected this set of legal and regulatory domains as they encompass a broad spectrum of approaches to handling ethical issues in research and represent a significant portion of the world's population. This way, we can highlight the benefits and disadvantages of varying solutions to regulating ethical issues.

We focused on ERBs that are publicly available to individual researchers and research groups employed by or affiliated with academic institutions. ERBs can be structured in very different ways.

We have included commercial ERBs as their financial interest in pursuing contracts with academic institutions inevitably leads to a certain level of accessibility. We explicitly omit internal governmental and military-specific ERBs, which US institutions heavily employ. As we argue in Sec. 4.2, we claim transparency and an open protocol approach to be crucial for such an institution. In this case, their necessary secrecy and the unavailability of non-military personnel means this is out of the scope of this paper. Also, military research is not usually submitted for review at conferences and journals in the CS community.

Please note that ERB structures are often designed for clinical trials and hence imply biomedical ethics, which in turn strongly influence oversight structures in CS. Therefore, examples of ERB structures that are solely responsible for clinical trials are given where no CS-specific structure was evident.

3.1 United States

In the United States, ethical approval is required for human subjects research, but it is not being regulated by a single, nationwide agency. Instead, it is governed by ERBs[3] that are present at most if not all, universities. These ERBs commonly create a set of forms that must be filed by the principal investigators (PIs) of research projects [64].

The content on these forms and the requirements to be met can vary widely between universities. For example, the University of Connecticut has a list of forms relevant to conducting research, which can be filled out [9]. In contrast, the University of Minnesota's ERB has compiled several long PDF documents outlining requirements for conducting human research, spearheaded by a 117-page "Investigator Manual" that contains numerous references to forms and other documents to be studied prior to research [42]. Lastly, several legal requirements surround issues of informed consent, and depending on the funding source for a specific research project, special requirements can be mandated. Most Federal Funding programs require an ERB approval for research, including Human Subjects [65].

3.2 Sweden

Research concerning people or biological material at Swedish research institutes must be approved by the Swedish Ethical Review Authority (SERA [80]). This authority is in charge of approving any research touching ethical questions. The process is highly regulated, both bureaucratically and legally. The application process works as follows: When research projects that (may) touch upon ethical issues are drafted, an application to SERA is mandatory. It has to contain a specially designed form with copies of all textual material of the research on human beings (e.g., for survey studies, a copy of the entire questionnaire must be attached to the application). After submission and payment of a mandatory fee, a decision will be made within 60 days and can result in an acceptance, rejection, or acceptance with minor or major changes.

A further requirement is that it is mandatory to request ethical approval before the project's actual start [79]. Additionally, if an approved project does not start within two years of the final acceptance by SERA, the waiver expires and has to be requested anew, again resulting in a fee to pay. These scheduling constraints, alongside the fact that all material that will be used during the (human) trials or surveys has to be provided before the project starts, require tight scheduling by the PI. After the project has been approved, no changes to any of the material are possible except by requesting a partial review of those changes. Legally, apart from being required to ask for approval by SERA, any violation of the ethical review rules can result in a fine or even imprisonment [58].

[3]They are usually referred to as Institutional Review Boards because they often also implement legal aspects.

3.3 Germany

Ethical governance in Germany is legally mandatory for the medical sector, specifically for research on drugs and medical devices. As research and education in Germany are the domain of the federal states, there are no competencies at the national level. In other fields, there is no legally mandated ethical governance.

Research concerning personal data is also increasingly being debated on whether mandatory ethics committees are necessary [88]. This is due to institutional pressure and progressing standardization within the Anglo-American scientific community. Currently, there are no legally binding requirements for handling personal research data beyond general provisions, e.g., the GDPR and its national equivalent implementation. Existing non-binding guidelines are set by funding agencies such as the German Research Foundation ("Deutsche Forschungsgemeinschaft", DFG) and professional societies like the German Informatics Society ("Gesellschaft für Informatik", GI). Notably, though, while the German Sociological Association does have its own ERB, the GI lacks such a body. Thus, especially when it comes to handling personal data, the EU-wide GDPR rules are the only legal framework that applies to such research in Germany [18].

3.4 European Union

The European Union, as a supranational organization, forms a governance layer atop national governance for its member states. Apart from nation-level requirements, distinct EU-led legal frameworks are emerging that further regulate the handling of sensitive user data. Examples of this are the EU Machinery Regulation [19] and the EU AI-Act [20]. These legal frameworks explicitly exclude research, development, and prototyping and only apply after a product is placed on the market. Thus, this legislation does not hinder academic research.

Specific standard rules for RCR can be found in the grant applications for the European Research Council (ERC) funding schemes, which entail mandatory ethics self-assessment [31] and have the character of non-binding guidelines. This assessment covers topics such as human rights and protection of human beings, animal protection and welfare, data protection and privacy, health and safety, environmental protection, and artificial intelligence [31, p. 23]. Any ethical requirements are settled between the ERC and the grantee, but apart from applicable national, international, and EU law, there are no legally binding principles targeting research in particular. Solely the ERC "strongly encourages" grantees to follow the ethical principles laid out by institutions of the European Union.

3.5 China

China has recently pushed to strengthen its own ethics oversight of science and technology. This goes back to the case incident in which a scientist presented the first genome-edited embryos without the knowledge of the government [53]. An updated guideline was published in March 2022 [91], requiring institutions to establish ethics committees and review research concerning humans and animals. The guideline also recommends revised ethics training in scientific education. Notably, this guideline applies generally to science and technology, i.e., not only to biomedical research and clinical studies, and therefore goes beyond the requirements in, e.g., Germany. It is not yet possible to foresee the impact of this strengthening of ethical oversight [91].

3.6 India

In India, Ethics Committees (EC) are responsible for ethical oversight. These are located either at academic institutions or hospitals or operate independently. India was a hub for commercial clinical trials till the 2010s. There have been many irregularities, especially around a lack of adequately trained personnel, heavy workloads, and lack of administrative support. This gave the ECs the reputation of rubber-stamping research for the sake of protocol [81]. To address this and establish a centralized quality control, the Central Drugs Standard Control Organization (CDSCO) introduced the required registration of ECs in 2013. Before conducting any clinical study in India, it is mandatory to seek approval from an EC registered with the CDSCO.

According to Thatte and Padmaja [81], additional regulations brought additional transparency and accountability over the constitution and practices of ECs. However, continuous monitoring of ECs is still not standardized, and therefore it can not be conclusively assessed how successful these regulations were beyond initial reports [68, 81].

3.7 Independent ERBs

Ethical oversight is predominantly only mandatory in the biomedical and behavioral sciences, i.e., where the need to protect human participants is obvious. This kind of mandatory ethical oversight in the US is bound to federal funding [6]. As mentioned before, all research institutions that receive federal funding have to employ some kind of IRB oversight, whether academic or independent. Using ERBs is not mandatory if there is no federal funding. Still, in its place, industry and governments often employ non-binding ethical guidelines [3].

Especially, pharmaceutical companies prefer using independent ERBs for clinical trials, which are considered more efficient than their academic counterparts [44]. Measuring whether academic ERBs perform more or less efficiently than commercial ERBs is difficult. Independent IRBs claim comparable quality to academic IRBs, and in some countries, they are regulated by authorities like the FDA [66]. These ERBs are not transparent, so this claim is impossible to assess or confirm. The decision-making processes are opaque, and there is no public record of the decisions. It should also be noted that independent ERBs primarily focus on clinical trials, and that is where most of their expertise can be found. As this paper focuses on research ethics in CS, a comprehensive analysis of independent ERBs is out of scope. However, their methods can be revisited in the future when it comes to finding role models for concrete implementation or as soon as independent IRBs have a significant share in the ethical oversight of CS research.

3.8 Summary and Contextualization

Table 1 summarizes how the governance structures discussed in this section integrate ethical reviewing of research within their boundaries. However, this is only a small snapshot of reality. We cannot claim completeness here, as it is hardly practicable to examine all existing structures. The number of approaches to this topic shows one thing: researchers from different structures who submit their work to the same conferences have to follow different complex paths to ensure the ethical correctness of their work.

The Swedish system is highly centralized and legally mandatory, thus already taking place at the *ex-ante* step of any project. Researchers must outline their data collection procedures so that the SERA can identify potential ethical issues. As such, the mental load on the researchers is fairly low, and the expectations are clear. However, the prospect of paying a non-trivial fee before the assessment has even started might prevent researchers from engaging in projects that require handling of user data in the first place. In addition, this carries the risk that researchers face financial issues if the project is turned down.

Another concern is that the legal requirement for ethical reviews can have a chilling effect on research more broadly. Since ethical issues are governed by different principles than law, it can be difficult for individual researchers to decide whether they have to apply for ethical review. This legal barrier, in addition to the high costs, might make researchers more cautious as to what to research. In 2023, this has led to about 2,500 Swedish researchers signing an open letter to the Swedish parliament, demanding reform.[76]

The ethical clearance process in the U.S. offers different insights. While the requirements for ethical research and the criteria for what is ethically problematic are much more elaborate, this comes at a cost: Researchers have to study manuals with sometimes over a hundred pages in advance.

This shows the need for globally standardized procedures. A standardized set of requirements would reduce overhead, e.g., when switching institutions, as once learned, they apply globally. Also, the United States shows that decentralized procedures might impede the efficiency and speed of ethical reviews since the institutions cannot learn from each other and amass. The case mentioned in Sec. 1 might have been flagged earlier if the responsible ERB had access to other ERBs decisions on similar research projects.

Germany, lastly, does not require any ethical clearance for CS research if the research project is connected to the requirements of the EU and GDPR. While the European Union arguably has the most rigid privacy laws, such as the GDPR, few regulations directly target research. This means that as long as researchers do not break any laws, their research is potentially legal[4]. However, such a generic and broad legal regulation sets the outer boundaries of ethical research practices. While ERBs need to evaluate individual projects, a legal framework helps to capture unanticipated implications.

[4]Please note that what is legal is not necessarily ethical, while what is ethical is not necessarily legal.

Table 1: Overview of ERB systems discussed in Section 3, that apply to CS.

	Research Domain	Obligatory	Regulation	Body
United States	domain independent	✔	approval	individual ERBs
Europe	domain independent	✘*	guidelines	EU agency
Germany	(Bio)medical Sciences	✔	approval	individual ERBs
	Computer Sciences	✘*	guidelines	individual institutions
India	(Bio)medical Sciences	✘	guidelines	national agency
China	domain independent	✘	guidelines	national agency
Sweden	domain independent	✔	approval	national agency

*In individual cases, restricted by legislation, e.g., GDPR.

4 Rethinking the Status Quo Procedures in CS Research

In this section, we revisit the challenges of ethical review procedures in CS. Combining the lessons learned from Sec. 3 and related work, we outline the requirements for a novel framework design.

4.1 Challenges of ERB Procedures in CS Research

Considering our observations in Section 2, the key challenges of ethical review in CS research are timing, i.e., local requirements, not necessitating ex-ante review, and local capabilities, i.e., the reviewing body not able to identify ethical challenges. This results in a similar effect as a lack of an ex-ante review, as the shortcomings of the conducted ex-ante review only become apparent after the ex-post. In addition to the *timing of the ethical review*, looking at the governance and policy frame, we observe that:

- Ethical clearance in CS is mostly only required on an *if need be* basis, i.e., not binding.
- In several cases, the ethical review procedures remain opaque, i.e., not verifiable.

We consider these four issues when considering possible new procedures for the ethical review of CS research. Considering the lessons learned, it becomes evident that ERB procedures must be internationally standardized, easy to understand, and mandatory for every research project. Additionally, the ERB procedures must be embedded in established and battle-tested practices such as peer review and grant applications so that both sides of the ethical clearance process, reviewers and the researchers, rely on existing expectations (e.g., what happens during a review).

To enable a fair and constructive ethical review, we advocate for a standardized procedure taking place on the *ex-ante* level. This allows questions of research ethics to be discussed beforehand, i.e., after the development of the research design and before experiments are conducted. This way, consideration of research ethics and a common framework are an inherent part of the research process and no longer an afterthought.

Implementing a community-wide framework addressing these challenges would have benefits for the community as a whole, going beyond the benefits of more ethical research:

- When ethical reviews are conducted *ex-ante*, conferences and journals are relieved of this task. They can focus on submissions' epistemic and methodological merits instead of being concerned about ethical compliance.
- A common ethics framework (not prescriptive ethics) avoids disparity between researchers required to submit an ethics review at their institution or state and those not, ensuring a level playing field.
- Researchers who do not have access to an ERB at their institution or whose ERBs do not have the necessary scientific expertise for a review can profit from a common pool of knowledge, where all other ERBs store their successful submissions and reviews. This ensures that under-resourced institutions are not at a disadvantage.

At this point, our proposal can be argued to promote the global standardization of normative principles and favor Western concepts of ethics. However, we argue that such developments can be discovered, scrutinized, and, most importantly, balanced through the quantity and diversity of ERBs and the exchange of experience and data. To put it more bluntly, by no longer tying ethical review to local ERBs, researchers must follow a rough consensus of different ethical perspectives, not advantaging Western perspectives alone.

4.2 Requirements

In the following, we propose testable properties that a new ERB procedure at least has to comply with to mitigate the discussed issues mentioned above.

R1: Standardized procedures Project submissions, ERB management, and the decision process should be processed using unified and deterministic procedures. Each research project can be handled equally, even by different ERBs. In turn, concepts like load-sharing between ERBs, or using a double-blind approach where ERBs do not necessarily review submissions from their "own" institution, become feasible.

R2: Transparency The procedure and outcome of ERB submissions should be accessible to all other ERBs in a searchable and transparent way. On the one hand, a transparent procedure visible to the scrutiny of the research community promotes traceability. This increases the chance that faulty ERB decisions are noticed and therefore reduces the incentive of ERBs to be negligent. On the other hand, ERBs working on similar cases can profit from these insights and, therefore, contribute to a unified and fair decision process. The same applies to the researcher; successful cases of ethical reviewing could be open to other researchers after their consent and publication of the corresponding work/paper. If they are working on similar projects, they can inform themselves in advance on how to prepare their project correctly in case of ethical concerns.

R3: Impartiality Researchers should have access to an ERB while avoiding conflicts of interests. This way, biased project assessments can be reduced as much as possible. If any such conflict arises, consulting a second ERBshould be possible.

R4: Domain Knowledge To understand whether a research project poses ethical concerns, technical domain knowledge, or in the case of specific populations, introspective experience is often required. However, ERBs are rarely staffed by computer scientists [55]. This is further aggravated by the amount of new disciplines in CS. Therefore, in case of

misunderstandings, ERBs should gain access to already processed examples of research projects or specialists of the sub-discipline of the submitted project without having to rely solely on the possibly biased project authors.

R5: Ex Ante Reviewing The fact that more and more conferences have recently established their own research ethics committees (REC) [12, 14] shows the importance of this topic. However, these RECs would come into play *ex-post*, i.e., when the research has already been conducted and potential harm has already been done. Therefore, any system fit to replace the current practices must require ethical reviewing to be conducted ex-ante. We do not intend a compromise between ex-post and ex-ante reviewing: We argue that potential harm can only be averted if thought on it is given in advance, i.e., the experimental design is reviewed ex-ante.

5 FEB Federation Framework

From this section onwards, we refer to boards that perform ethical reviewing of research projects, similar to ERBs, as Federated Review Boards (FEB). We describe a framework for organizing FEBs in an association that we refer to as *FEB Federation* throughout this work. This association is represented by a centralized entity in charge of managing project submissions, automatically assigning submissions to suitable FEBs, and processing the data resulting from each decision process. In Sec 5.1, we propose the roles for such an FEB Federation and give a high-level overview of the procedure in Sec 5.2.

Precondition: Registration. Each research institution, e.g., a University, that wants its researchers to profit from the proposed FEB procedure must, in turn, contribute to it, e.g., by registering its own FEB to the FEB Federation or by cooperating with another institution, that already contributed. Only then are researchers from this institution authorized to submit their work to the FEB Federation. This ensures that interaction with the federated structure is reciprocal. Recall that in R4 (Sec. 4.2), we assume the FEB participants not to be experts in all *specific* sub-fields of CS or human-related sciences. Therefore, during registration, the FEB can specify domain tags that indicate the general research disciplines they can consult at their hosting University/Research Institution, e.g., a robust research track in Web Security, tight collaboration with a medical faculty, specific expertise and self-reflective experience with specific communities, etc.

5.1 Framework Roles

The FEB Federation we propose requires different roles that must participate or be present during the procedure, described in the following. The interplay of all roles during the procedure split up into four phases is shown in Figure 1.

A **submitter** is an individual or a group of researchers associated with a research institution who work(s) on a project with possible ethical implications. They aim to submit their project idea to an FEB to get feedback and attestation on its ethical correctness, which they can include when submitting their paper to a conference.

An institutional review board (**FEB**) supports researchers in matters of ethically correct research and is in charge of reviewing submitted projects. Within the institution that hosts the FEB, it has access to research domains (university faculties, institutes), which it can consult if questions occur regarding domain-specific projects, e.g., Web Security, Machine Learning, or User Privacy research.

The **FEB pool** refers to all pre-registered FEBs. Submitter who are associated with an FEB are authorized to submit their projects to it. During registration, each FEB specifies research domain-specific expertise it has access to in the form of domain tags, which are then associated by the pool with this FEB. Since the FEB responsible for the submission is selected randomly from the FEB pool, this minimizes bias due to local power dynamics.

The **pool control** is a centralized governance entity that maintains the FEB pool. It handles FEB registration, submissions, and the assignment of a submitted project to an FEB. In addition, it stores information about each finished FEB decision and provides access to this data to authorized entities. This entity could be a professional organization similar to the ACM or, e.g., the CHI community. While the specific implementation is out of scope for this paper, we discuss a potential governance structure in Sec. 6.4.

The **meta form** is a pre-defined form used to describe the [5] research project as part of the submission input. It contains both a standardized questionnaire to capture ethical considerations at a glance and meta-information regarding the project (such as the domain above tags) in machine-readable form so that the submitter can be matched to an FEB quickly. This meta form is filled out by the submitter and sent to the pool control. Note that in many cases, the submitter must draft a project plan in advance anyway, e.g., for funding.

A submitted work that has been reviewed by an FEB and passed successfully the procedure should be attested in an unforgeable way, e.g., using a *digital certificate* containing the signatures of the FEB and the pool control and targeting the meta form. The resulting attested meta form is denoted as $\mathbf{meta}^{cert}$. This method allows the traceability of any later modifications to the research project that were not reviewed by the FEBs[6]. This approach is similar to the badge attached to the paper, which has undergone the Artifact Evaluation process of conferences [1].

5.2 High-Level Overview

The idea is to provide a platform where each participating FEB is registered in advance and, therefore, is part of the FEB pool. Instead

[5] Elaborating the exact content of the meta form is out of scope in our work. However, examples from, e.g., the US FEBs as described in Sec 3.1 can serve as role models.

[6] A possible solution is the use of techniques like the Electronic Signature Attestation [35]. Technically, such a record could also be established in a non-technological manner, e.g., by notarizing a document. However, elaborating a suitable solution is out of scope for the presented work.

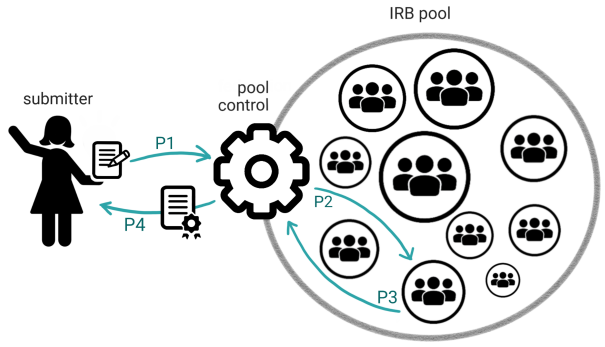

Figure 1: Overview of the review process within the FEB Federation

of manually contacting their responsible FEB, submitter submits the meta form to the pool control, containing a description of their intended project, possible conflicting FEBs, and domain tags that match their project's research discipline.

The pool control then selects a suitable FEB at random from all registered FEBs in the pool. To ensure the FEB can access the domain knowledge of the submitted meta form in case questions arise, it pre-filters the FEBs available for the selection by considering the domain tags of both, the meta form and all FEB s in the pool, except those who have been declared as conflicted. The FEB selected this way is then in charge of reviewing the submission. From here on, the FEB directly communicates with the submitter if any questions occur or adjustments to the intended experiments are needed.

The decision is then reported by the FEB back to the pool control. They can finally certify the meta form and send a $\textbf{meta}^{cert}$ to submitter. With their consent, e.g., after a specific time or a successful publication, the pool control adds the data of this review process to a joint database, this way providing the data to other FEBs and/or researchers.

5.3 Phases of the Procedure

In the following, we introduce our procedure as shown in Figure 1, split into different phases. Since the assignment of a suitable FEB and the certification afterward can be automated, there is no overhead for the submitters in addition to the current procedure. As with most conferences in CS, the review process itself is carried out in a double-blind manner using a communication system similar to the rebuttal process.

P1: Submission. At the beginning, the submitter prepares the meta form for their research project, as required by the FEB Federation. Due to their extensive nature, the standardized questions on the form already capture many common pitfalls (such as, "Do you process user data?"). As such, the submitter can fine-tune their project design before the application is sent out. The submitter submits the meta form to the pool control who accepts it if the hosting institution of at least one author registered an FEB to the FEB Federation. In case of conflicts, the submitter can declare any conflicting FEB within the pool.

P2: Assignment. Based on the information in the meta form, the pool control selects a suitable FEB at random that is then in charge of reviewing this submission. Identifying the exact criteria for this selection is outside the scope of this work since they may depend on additional factors. However, the domain tags of the possible FEBs should match the submission's domain tags. This way, it is ensured that the selected FEB has access to domain-specific support at their hosting institution if needed.

P3: Reviewing. The submission will then be reviewed by the FEB chosen by pool control based on the meta form. This allows an FEB to be selected that has the necessary knowledge for the submitted project and is, therefore, more likely to be able to assess the ethical considerations. To expand their competence, the FEB has two supportive possibilities: It can access a shared database of the FEB Federation where all decisions of the previous cases are stored. This way, they can determine how others have mitigated similar problems before greenlighting a project. Suppose the database lacks useful information from similar previous projects. In that case, the FEB can consult further experts in this domain at their hosting university, which is possible due to the matching domain tags. In addition, the FEB can directly communicate to the submitter if further information or adjustments to the project proposal are required to ensure ethically correct research. The submitter, in turn, has to respond with plans on how to mitigate identified problems. This exchange continues until all ethical issues have been addressed.

P4: Decision. Finally, the FEB confirms to the pool control that the specified project can commence. The submitter receives a $\textbf{meta}^{cert}$ from pool control confirming the ethical clearance and any identifier attached. Upon submitting the corresponding project paper to any conference or journal, the submitter can attach the $\textbf{meta}^{cert}$ to their paper. After the publication of the paper, with the author's consent, the DOIs of the publications are sent back to pool control so that they can be connected to the original meta form and include the project in the shared database of the FEB Federation.

Follow-up submissions. Later, another similar project is submitted for ethical review to the FEB Federation. Considering the matching domain tags, the submitter of the new project consults the database for publicly available data of any similar previous projects. This way, the submitter can understand the points they must consider before submitting. The FEB selected by the pool control has access to these additional resources as well: they can consult the data from the previous project and read through both the initial meta form sent, the changes made and the final publications to draw on these for a better and quicker judgment. The more submissions are processed this way, the larger this database grows and the more helpful the data is for the FEB and submitter to draw on.

5.4 Addressing the Requirements

We have identified five requirements for any replacement to the current ethical review practices must include: standardization, transparency, impartiality, domain knowledge, and reviews at the *ex-ante* level. Here, we discuss how the proposed FEB Federation addresses these.

- **Standardized procedures (R1)** aims to reduce the cognitive load for researchers and ethical reviewers alike by making the

ethical approval process more accessible to navigate. Our framework achieves this by centralizing the formal processes in a single entity: the pool control. Its role is to implement a standard protocol for the procedure of ethical reviewing, i.e., it affects *how* the decision should be made, but not *what* decision should be made. In other words, in our work, we propose standardizing the submission- and approval process but do not dictate how the actual decision will come about. The participating FEBs remain free in how they rule.

- Our proposed FEB Federation provides a high amount of **transparency (R2)**. Since the applications and decisions will be made public after the publication of the resulting papers, this creates trust among both reviewers and submitters and the (research) community in general. Reviewers will not suddenly craft an unprecedented decision, and both submitters and reviewers benefit from the increasing access to knowledge.
- Our procedure is highly **impartial (R3)**. Since the reviews are performed via a double-blind peer review procedure and FEBs are randomly allocated via pool control, the chances that personal conflicts might bias a decision are lowered. And even if a personal connection results in a false-positive or false-negative decision, this becomes public due to the transparency our framework facilitates. Here again, network effects are leveraged to ensure fair play by all parties.
- Our proposal allows the inclusion of a vast array of **domain knowledge (R4)** into the ethical review procedure. While each University generally only has one FEB with limited domain knowledge, our framework increases the chances that researchers will be able to submit their project proposal to a fitting FEB, as each new FEB brings additional domain knowledge into the FEB pool, or expands the existing.
- Researcher and publication venues only benefit from the FEB Federation if the submissions are accompanied by a $\textbf{meta}^{cert}$, such that the procedure ensures that ethical reviewing is always conducted **ex-ante (R5)**. This way, researchers can rest assured that their papers are unlikely to be rejected on ethical grounds. As the public knows that the FEB Federation serves as a strong disincentive against conducting unethical research, this also relieves conferences and journals of the current pressure to double-check submissions on ethical grounds if the submitted papers are marked with a valid $\textbf{meta}^{cert}$.

In summary, our proposed FEB Federation provides an adaptable system that evades local power dynamics, makes community resources more widely accessible, and ensures specific competence to be present during the ethical review. Furthermore, our transparency requirements enable a general 'just culture' approach of constant learning and refinement, helping the community grow and adapt to newly found circumstances, which we cannot anticipate yet[7].

6 Discussion

The following section is dedicated to ideas on how to reach the acceptance of the proposed FEB Federation within the research community, confronting some limitations and concerns and proposing some possible first steps to be done toward a unified and fair procedure of ethical reviewing.

6.1 Increasing Acceptance in the Research Community

Suppose such a new framework is to become an effective way to ensure safe and sound ethical reviews of research projects. In that case, it is paramount to ensure the affected actors do not attempt to subvert it or outright refuse to participate. Formally, the FEB Federation must add as few novel processes to the mix as possible. In other words, the FEB Federation must look and feel like institutions that researchers and organizations have come to accept. This section discusses strategies to mitigate the danger of subversion or refusal.

Institutionalism describes the process by which specific procedures become accepted in a given field and followed by organizations without the ambition to bypass them [25, 57]. The essential academic institutions that our proposal relates to are peer-reviewing and grant applications. These are institutions that institutionalism describes as being *isomorph*. That is, they are being followed by all actors involved unquestioned. Since such institutions are tried and tested, they provide safety and trust and come with a set of known expectations.

For example, when a researcher submits a paper for publication, they expect other researchers to review it and that they have to modify the paper before it will be accepted, e.g., as a major revision. Conversely, grant applications require the applicant to conduct thorough research beforehand and discuss as many potential pitfalls the project might encounter as possible. Researchers writing grant applications expect that their application will be successful if their prior research was complete (as in: addresses all problems to be expected). This institutionalization process was later referred to as providing "cultural schemata" [24]. By tapping into these two known procedures, we can ensure that the actors will accept our proposed FEB Federation and thus ensure that ethical review is being conducted faithfully. On the other hand, if this framework would introduce many new procedures none of the involved actors are comfortable with, this results in a risk of "decoupling", a process in which the formal requirements of a procedure would still be followed, but internally the expected safeguards are bypassed[8]. By framing ethical reviews in terms of peer review and grant applications, we ensure that the researchers involved know what to expect and understand the necessity of ethical review.

6.1.1 Additional advantages of the FEB Federation. On a more general level, there are further benefits that can help facilitate acceptance of the introduced FEB Federation. The possibly most significant incentive for researchers to participate is receiving competent, ethical guidance by FEBs. This means they are not left alone in making their research safe.

Universities likewise have an incentive to join the FEB Federation since they do not only provide services to the FEB pool but also benefit from other Universities' FEBs and, therefore, domain knowledge. Participation will decrease the workload for the University's own FEB: Instead of having to review each research proposal

[7] Predicting the future is notoriously difficult; Hence, it is instrumental to not fail at an attempt to build a system that can handle all foreseeable issues, but instead creating a system that can adapt, even in the face of unforeseeable challenges.

[8] Granovetter describes in [40] a detailed example of managers in a company going to extreme lengths to avoid critical, internal audits by the central management.

from their University regardless of their domain knowledge, they can specialize in one or a few domains, which increases the FEBs speed as well as the quality of its decisions.

Furthermore, the data itself can be used to analyze ethical issues in research on a broader level. We identify precisely three types of actors that will benefit from the data the FEB Federation generates: (a) researchers themselves can reduce uncertainty before submitting a project for review; (b) the reviewers can draw on past decisions to inform themselves; and (c) disciplines such as Science of Science [38] will receive a new data source to learn from.

Lastly, there are even benefits for Universities without an FEB to create one. An FEB can be incorporated into the existing syllabus to teach ethical practices, as it already has been proposed [34, 85]. This will raise the awareness of CS students, i.e., future researchers, so ethical problems are less likely to arise during review. Training CS students in ethical procedures can proactively diminish incidents like those discussed initially.

6.1.2 The Role of Rough Consensus. As noted in Section 2.1.2, no *one* set of ethics will hold across all cultures and disciplines. Hence, for FEB Federation, we argue that a rough consensus approach to ethics across FEBs as introduced in Section 2.1.3 is the ideal method to converge to a global state that is not the same, but *acceptable* for all involved.

However, implementing such a model will require engagement between researchers and FEBs across institutions, i.e., throughout all institutions participating in FEB Federation. Hence, tying participation in the framework - and potentially governance of FEB Federation itself - to participation in a regular, e.g., annual meeting during which guidelines are established in a bottom-up approach would be a feasible option to encourage quick convergence, as well as discussion of current issues. We note here, though, that this concept of a "meeting" or "conference" is not necessarily aligned with the publication-driven nature of symposia in CS; As such, it might require some additional attunement within the community.

6.2 Limitations

One main limitation of our proposal is that it is based on a threshold assumption [39]. That is, as long as the FEB Federation does not cross a certain participation threshold, the first FEBs will likely have to deal with an increased workload. This could potentially create a free rider-effect, in which Universities attempt to remain out of the FEB Federation for as long as possible to benefit afterward. In our proposal, we address this through the tit-for-tat [61] based requirement for institutions to contribute to be part of the framework.

This incentive structure becomes increasingly effective as more and more conferences and journals endorse the FEB Federation, reducing the costs of joining the proposed FEB Federation. Specifically, this applies to conferences and journals that have begrudgingly taken it upon them to vet research *ex-post*. Currently, in their reviewing procedure, many conferences in CS offer an input field for the reviewers to flag papers where they think there might be unresolved ethical issues. Those venues are incentivized to promote a centralized *ex-ante* approach where they can eliminate this additional responsibility and only focus on the quality of the submitted research.

A second limitation is that our framework still faces the problem of deciding whether ethical review is necessary in the first place. In the aforementioned case, the researchers allegedly were not aware that they were conducting ethically problematic research. This points to the fact that ethical review is not just a matter of unified and standardized procedure but also of culture. We argue that our approach allows a broader base of perspectives to be used in ethical review. Specifically, due to the possibility to leverage the experience and specific expertise of distributed FEBs, it becomes "easier" to use a "just submit for review" stance. In such a setup, even domain-specific topics that commonly do not require in-depth reviews will be placed with FEBs with specific experience in the matter, able to provide accelerated assessment of clear-cut no-review cases.

This mitigates not only the issue of overlooking cases where ethical review is necessary but also the case where an FEB not experienced in a specific subject performs a lengthy - often longer due to additional time needed for due diligence and to resolve misunderstandings - review of a project that would not require ethical review. An FEB with domain-specific knowledge could accelerate this phase.

A third limitation concerns the structure of the required meta form. Different types of research require different information to be provided to the FEB Federation. While medical research needs to provide information on, e.g., whether ionizing radiation will be used on research subjects to help the reviewers weigh the risks and benefits of the research (beneficence), this is not relevant for CS research. Therefore, the meta form must be an iterative process, subject to constant changes and even conflicts among the FEBs. As stated above, specific cultural settings may make some aspects more or less relevant to the participating FEBs. This could lead to conflicts about what information the FEB Federation requires from its submitters. However, as stated previously, such forms are already used by present institutions that can serve as role models.

6.3 Concerns

As our proposal requires fundamental structural changes, it also comes with some concerns. Contrary to limitations, i.e., points that could be improved and have to be acknowledged, we also see concerns, i.e., aspects of our proposal that create new *risks*. Identifying and addressing all of these requires the involvement of different actors from within the CS research community. In the following, we discuss the concerns that we identified.

First, creating an extensive *public database* naturally raises concerns on several levels. The database necessarily contains private information by the submitter, subject to legal frameworks such as the GDPR. It is, therefore, indicated that the database should contain one public and one private part. The public part could include the initial applications, the decisions, and the resulting papers and be accessible as soon as the related project is published. The private part could include only the information necessary to facilitate communication between the FEB and the submitter. If any issues or conflicts arise later, it could be made available to conflicting parties, e.g., another FEB or publishing conference, to trace the past decision process[9].

[9]This is only a suggestion from our side. The precise design is only possible with the involvement of the research and publishing community.

A second concern is *scooping*. Scooping is the term for having someone else claim priority on a research idea, usually through publishing, one has worked on independently [51]. It can happen by chance that someone is working on it simultaneously or that a competing researcher adopts one's idea by learning from leaked information before publication. To prevent information leaks from the FEB pool, the implementation must ensure that information on papers under submission remains confidential until the consent of the submitter, e.g., after the paper's publication date. Utilizing techniques like time-stamping the submitted research ideas could allow the authors to prove scooping afterward. However, this issue is similar to the peer-review process in research conferences, and their policies can serve as a role model for ex-ante ethical reviewing.

A third concern regards *faulty decisions*. As with any system, our proposed FEB Federation may suffer from faulty decisions, specifically Type I and Type II errors, i.e., projects being approved even though they are ethically concerning and projects not being permitted despite being ethically viable.

On the surface, Type I errors can be addressed by implementing a mechanism for retraction akin to that of journal articles[10]. In 2019, the reference manager Zotero, for example, included a feature that marks retracted papers in red and advises against citing them[11]. Such a retraction mechanism could work fairly similarly to journal retractions, with one crucial difference:while papers can be retracted due to a variety of reasons - including ethical problems identified *ex-post* - only, retractions due to ethical considerations are of importance to the FEB Federation.

For Type II errors, we argue that the proposed appeal process should be able to resolve cases, especially given that it allows the inclusion of increasingly more FEBs. This is also important if a research project is proposed for which no rough consensus can be obtained within the community. In that case, we argue that involving growing parts of the community in the discussion is paramount.

More crucially, for Type I and II errors alike, it is imperative that the system 'learns' from these mistakes. In these cases, the whole ethical review process must be flagged to inform both researchers and FEBs alike that there may have been noteworthy circumstances in the ethical approval of the project. There, the public database (see the transparency requirement discussed earlier) enables FEBs to *learn* from (others') prior mistakes, essentially building a living body of self-improving guidelines ensuring adaptability to ever-changing circumstances. This approach should be supported and informed by an overall 'just culture' approach [21] and is similar to, e.g., how aviation safety works, and has also been proposed for, e.g., IT security best practices [32].

A last concern regards *already established systems* of ethics, especially where some or all parts of the ethical review process are legally mandated, as in the case of Sweden (see Section 3). As the proposed FEB Federation is organized on a transnational level, it will inevitably clash with any national regulation that mandates certain forms of ethical review. Under current Swedish law, the SERA would have no choice but to disregard the authority of pool control over the ethical vetting process. Any researcher seeking vetting by the FEB Federation must submit two applications.

This will put researchers in countries where there are strict national regulations on ethical review processes at a disadvantage. However, we have two expectations in this regard. First, as both authorities review the same project, we expect the overhead from writing two applications to be manageable. But, second and more importantly, we expect that the FEB Federation will prove its worth to the legislative bodies of such countries so that they eventually agree to abandon national regulation and instead join the FEB Federation.

6.4 Implementing the FEB Framework

As with all ideas, the main challenge of this work is *actually* implementing it, i.e., making it happen. We argue that the best way forward is to document our proposal—which we do in this paper—and discuss it in a community already actively working on improving the ethics processes in CS. The CHI/CSCW community is currently at the forefront of ethical discussions within CS, ranging from matters of positionality via informed consent questions to questions of the transitive impact of research via its outputs.

While questions of procedure can be addressed in a conceptual framework, a fundamental part of the proposed FEB Federation is one of trust in the pool control. Implementing a proper governance structure to coordinate the entire framework is crucial to fostering trust by the community.

A specific governance model for the pool control is out of scope for this paper as it moves too far into the realm of implementation. We note that the governance structure of such a body must follow a bottom-up self-governance structure by the community. This must remain independent of any established institutionalized entity with its potentially own agenda to ensure impartiality. Conveying the responsibility for the FEB Federation to an existing organization runs the risk of decreasing the global level of trust, as existing biases and power imbalances will be transferred from the organization to the pool control. A novel governance body is likely to make the FEB Federation as acceptable as possible. Given that our FEB Federation essentially can scale from a minimal set of participating FEBs - worst case at least two - we argue that the only thing we have to do is finding a second institution (or, ideally, more) willing to participate in the proposed FEB Federation. Because everything is better than the situation we find ourselves in now, we can take it from there, following the concept of adaptability.

At this point, we cannot and do not want to offer a ready-made and perfect proposal for implementation. This requires the greater involvement of the research community and the relevant stakeholders. Furthermore, we believe the search for a perfect solution at the beginning is prone to falling victim to a Nirvana Fallacy[12]. To avoid this, we refer to the following two principles, known from the domains of Agile Software Development and System Administration.

The "Incremental Deployability" approach refers to a practice used in Agile Software Development to gradually and iteratively deploy new features or updates in production environments. For

[10] A retraction mechanism for faulty papers is already institutionalized.
[11] https://www.zotero.org/blog/retracted-item-notifications/

[12] Also known as "Perfect Solution Fallacy", coined by Demsetz [22]. It refers to pursuing an idealized, perfect solution, ultimately leading to potentially imperfect but practicable solutions not even being considered.

example, a new feature is divided into stages that are applied incrementally and, if possible, only to a subset of the system. There exists another concept that makes an excellent synergy: The Small Badges Principle, as described by Limoncelli [54], refers to a strategy which does not attempt to build the "perfect system", but rather to build a small system first and develop it further. It also considers social aspects like stress and morale. Inevitably, the implementation of a new system as the FEB Federation proposed by this work will initially involve additional work for the involved subgroup. However, we propose to combine both principles to create an agile ethics reviewing procedure for the transition that is centered on user feedback and causes fewer overloads for researchers, their institutions, and publishers.

7 Concluding Remarks

In this paper, we propose a FEB Federation to fundamentally improve the ethical review process in CS in a collaborative and transparent manner. Starting from exploring challenges encountered in the status quo, we derive design requirements for ethical review for human-affecting studies, using the Security and Privacy domain as an example.

We argue that current procedures are insufficient and sometimes even arbitrary, and disparities among them can result in ethically questionable approaches being overseen or misjudged. Individual researchers may even be disadvantaged depending on their affiliation or financial possibilities. Furthermore, we find challenges in local and community-wide power structures and limitations in whether ERBs and researchers can be reasonably expected, i.e., can be expected to be qualified to assess the full scope of ethical considerations, primarily when *subjective* harm must be assessed when, e.g., marginalized communities are being surveyed.

From these limitations, in Sec. 4.2, we derive five requirements $R1-R5$ that must be met by any new process for reviewing research ethics. Leveraging these requirements, we introduce the concept of a FEB Federation for the federation of ethical review resources throughout the community. Contrary to the current state of the art, i.e., *reactive* ex-post review leading to, e.g., paper retractions while *harm has already been done*, we propose a federated community effort. Our proposal also approaches the large disparities of current procedures of ethics reviewing in research due to sociocultural differences in assessing what *is* ethical, focusing on continuous iteration and improvement. Adopting a transparent and non-local procedure ensures a continuous alignment with a community-established standard of 'rough consensus.' Finally, we show how our FEB Federation can be embedded in the current research process and discuss its boundaries and possibilities to increase the acceptance of our proposal within the research community.

In summary, we call on the research community to join forces and propose a unified procedure for ethical reviews of CS research projects in this work. We are aware that this endeavor is an enormous challenge for the research community, especially since rules need to be agreed to by all involved institutions. As the Internet and digital technology, in general, have become integral to how societies work worldwide, they influence humans and vice versa. This results in implications of CS work regularly transcending the matter of pure bits. Hence, it is imperative to tackle the ethical challenges of our collective work. Not only to act ethically but also to foster trust in CS research and enable researchers to conduct research for the benefit of society.

Acknowledgments

We greatly appreciate the valuable input from our reviewers and the participants at NSPW'24, which has helped to enhance our work through insightful discussions and constructive feedback. S.G. was supported by funding from the Baden-Württemberg Foundation (BWS) under the Autonomous Robotics 2023-2026 call.

References

[1] 2015. Artifact evaluation. In *2015 IEEE 3rd Working Conference on Software Visualization (VISSOFT)*. vii–viii. https://doi.org/10.1109/VISSOFT.2015.7332408
[2] Lura Abbott and Christine Grady. 2011. A Systematic Review of the Empirical Literature Evaluating IRBs: What We Know and What We Still Need to Learn. *Journal of Empirical Research on Human Research Ethics* 6, 1 (March 2011), 3–19. https://doi.org/10.1525/jer.2011.6.1.3
[3] AlgorithmWatch. 2020. AI Ethics Guidelines Global Inventory. https://inventory.algorithmwatch.org/.
[4] Ronald E Anderson. 1992. ACM code of ethics and professional conduct. *Commun. ACM* 35, 5 (1992), 94–99.
[5] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The Menlo report. *IEEE Security & Privacy* 10, 2 (2012), 71–75.
[6] Michael Battaglia, N Sampling, and PJ Lavrakas. 2008. *Institutional Review Board (IRB)*. Sage Publications, Inc., 2455 Teller Road, Thousand Oaks California 91320 United States of America. https://doi.org/10.4135/9781412963947.n225
[7] Tom L Beauchamp and James F Childress. 2019. Principles of Biomedical Ethics. (2019).
[8] Rasika Bhalerao, Vaughn Hamilton, Allison McDonald, Elissa M. Redmiles, and Angelika Strohmayer. 2022. Ethical Practices for Security Research with At-Risk Populations. In *2022 IEEE European Symposium on Security and Privacy Workshops*. 546–553. https://doi.org/10.1109/EuroSPW55150.2022.00065
[9] Douglas Bradway. 2014. IRB Forms & Templates | Office of the Vice President for Research.
[10] Elizabeth Buchanan, John Aycock, Scott Dexter, David Dittrich, and Erin Hvizdak. 2011. Computer science security research and human subjects: emerging considerations for research ethics boards. *Journal of empirical research on human research ethics: JERHRE* 6, 2 (June 2011), 71–83. https://doi.org/10.1525/jer.2011.6.2.71
[11] Elizabeth A. Buchanan and Michael Zimmer. 2023. Internet Research Ethics. In *The Stanford Encyclopedia of Philosophy* (Winter 2023 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University.
[12] Joseph Calandrino and Carmela Troncoso. 2023. Message from the USENIX Security '23 Program Co-Chair. https://www.usenix.org/sites/default/files/sec23_message.pdf
[13] Augustin Chaintreau, Pan Hui, Jon Crowcroft, Christophe Diot, Richard Gass, and James Scott. 2007. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Transactions on Mobile Computing* 6, 6 (2007), 606–620.
[14] USENIX Security '23 Program Chair. 2023. Call for papers, Research Ethics Commitree. https://www.ieee-security.org/TC/SP2023/cfpapers.html
[15] Markus Christen, Bert Gordijn, and Michele Loi. 2020. *The ethics of cybersecurity*. Springer Nature.
[16] Yi Ting Chua, Simon Parkin, Matthew Edwards, Daniela Oliveira, Stefan Schiffner, Gareth Tyson, and Alice Hutchings. 2019. Identifying unintended harms of cybersecurity countermeasures. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 1–15.
[17] Usenix'23 Programm Committee. 2023. Statement from the USENIX Security '23 Program Committee. https://www.usenix.org/system/files/sec23winter-prepub-226-kirchner_statement.pdf
[18] Marc Cornock. 2018. General Data Protection Regulation (GDPR) and implications for research. *Maturitas* 111 (2018), A1–A2.
[19] Council of European Union. 2024. Regulation - 2023/1230 - EN - EUR-Lex. https://eur-lex.europa.eu/eli/reg/2023/1230/oj. Doc ID: 32023R1230 Doc Sector: 3 Doc Title: Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery and repealing Directive 2006/42/EC of the European Parliament and of the Council and Council Directive 73/361/EEC..
[20] Council of European Union. 2024. Regulation - EU - 2024/1689 - EN - EUR-Lex. https://eur-lex.europa.eu/eli/reg/2024/1689/oj. Doc ID: 3202/4R1689 Doc Sector: 3 Doc Title: Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonized rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU,

(EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act)Text with EEA relevance. Doc Type: R Usr_lan: en.
[21] Sidney WA Dekker and Hugh Breakey. 2016. 'Just culture:'Improving safety by achieving substantive, procedural and restorative justice. *Safety science* 85 (2016), 187–193.
[22] Harold Demsetz. 1969. Information and efficiency: another viewpoint. *The journal of law and economics* 12, 1 (1969), 1–22.
[23] Department of Computer Science and Engineering. 2021. https://cse.umn.edu/cs/news/paper-accepted-ieee-symposium-security-and-privacy-2021. archived at https://web.archive.org/web/20210421135935/https://cse.umn.edu/cs/news/paper-accepted-ieee-symposium-security-and-privacy-2021.
[24] Paul DiMaggio. 1997. Culture and Cognition. *Annual Review of Sociology* 23, 1 (1997), 263–287. https://doi.org/10.1146/annurev.soc.23.1.263
[25] Paul DiMaggio and Walter Powell. 1983. 'The Iron Cage Revisited: Isomorphism in Organizational Fields'. *American Sociological Review* 48 (April 1983), 147–160.
[26] Verena Distler, Matthias Fassl, Hana Habib, Katharina Krombholz, Gabriele Lenzini, Carine Lallemand, Lorrie Faith Cranor, and Vincent Koenig. 2021. A systematic literature review of empirical methods and risk representation in usable privacy and security research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 6 (2021), 1–50.
[27] David Dittrich, Katherine Carpenter, and Manish Karir. 2014. An ethical examination of the Internet Census 2012 dataset: A Menlo report case study. In *2014 IEEE International Symposium on Ethics in Science, Technology and Engineering*. IEEE, 1–5.
[28] Roya Ensafi, Philipp Winter, Abdullah Mueen, and Jedidiah R Crandall. 2015. Analyzing the great firewall of china over space and time. *Proc. Priv. Enhancing Technol.* 2015, 1 (2015), 61–76.
[29] Paul Ernest. 2018. The ethics of mathematics: Is mathematics harmful? *The philosophy of mathematics education today* (2018), 187–216.
[30] Linderud et. al. 2023. Ethics concerns about paper accepted for publication. https://hackmd.io/s/BJGs6Tfiw
[31] European Commission. 2021. Horizon 2020: Online Manual. https://ec.europa.eu/research/participants/docs/h2020-funding-guide/cross-cutting-issues/ethics_en.htm.
[32] Tobias Fiebig. 2020. How to stop crashing more than twice: A Clean-Slate Governance Approach to IT Security. In *2020 European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 67–74.
[33] Tobias Fiebig. 2023. Crisis, Ethics, Reliability & a measurement. network: Reflections on Active Network Measurements in Academia. In *Proceedings of the Applied Networking Research Workshop*. 44–50.
[34] Casey Fiesler, Mikhaila Friske, Natalie Garrett, Felix Muzny, Jessie J Smith, and Jason Zietz. 2021. Integrating Ethics into Introductory Programming Classes. In *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*. 1027–1033.
[35] Benjamin David Follis. 2014. Attestation for electronic signatures.
[36] Romain Fontugne, Ksenia Ermoshina, and Emile Aben. 2020. The Internet in Crimea: a case study on routing interregnum. In *2020 IFIP Networking Conference (Networking)*. IEEE, 809–814.
[37] National Commission for the Protection of Human Subjects of Biomedical and US Behavioral Research. 1978. *The Belmont report: ethical principles and guidelines for the protection of human subjects of research*. Vol. 2. The Commission.
[38] Santo Fortunato, Carl T. Bergstrom, Katy Börner, James A. Evans, Dirk Helbing, Staša Milojević, Alexander M. Petersen, Filippo Radicchi, Roberta Sinatra, Brian Uzzi, Alessandro Vespignani, Ludo Waltman, Dashun Wang, and Albert-László Barabási. 2018. Science of Science. *Science* 359, 6379 (March 2018). https://doi.org/10.1126/science.aao0185
[39] Mark Granovetter. 1978. Threshold Models of Collective Behavior. *The American Journal of Sociology* 83, 6 (May 1978), 1420–1443.
[40] Mark Granovetter. 1985. Economic Action and Social Structure: The Problem of Embeddedness. *Amer. J. Sociology* (Nov. 1985).
[41] Florian Hantke, Sebastian Roth, Rafael Mrowczynski, Christine Utz, and Ben Stock. 2024. Where Are the Red Lines? Towards Ethical Server-Side Scans in Security and Privacy Research. In *IEEE Symposium on Security and Privacy*.
[42] Hsantiag. 2017. HRP-103, Investigators Manual. https://research.umn.edu/units/irb/toolkit-library/manuals
[43] IEEE S&P 21 Program Committee. 2021. IEEE S&P'21 Program Committee Statement Regarding The "Hypocrite Commits" Paper. https://www.ieee-security.org/TC/SP2021/downloads/2021_PC_Statement.pdf
[44] Sheila Kaplan. 2016. In clinical trials, for-profit review boards are taking over for hospitals. Should they? https://www.statnews.com/2016/07/06/institutional-review-boards-commercial-irbs/
[45] Mannat Kaur, Michel van Eeten, Marijn Janssen, Kevin Borgolte, and Tobias Fiebig. 2021. Human factors in security research: Lessons learned from 2008-2018. *arXiv preprint arXiv:2103.13287* (2021).
[46] Tobias Keber. 2021. Digital Ethics by Process? Technical conflicts and policy ethics committees in Europe. *Informatio. Revista del Instituto de Información de la Facultad de Información y Comunicación* 26, 1 (2021), 216–229.
[47] Robin Kirchner, Jonas Möller, Marius Musch, David Klein, Konrad Rieck, and Martin Johns. 2024. Dancer in the Dark: Synthesizing and Evaluating Polyglots for Blind Cross-Site Scripting. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association.
[48] Robert Klitzman. 2011. How Local IRBs View Central IRBs in the US. *BMC Medical Ethics* 12, 1 (Dec. 2011), 13. https://doi.org/10.1186/1472-6939-12-13
[49] Tadayoshi Kohno, Yasemin Acar, and Wulf Loh. 2023. Ethical frameworks and computer security trolley problems: foundations for conversations. In *Proceedings of the 32nd USENIX Conference on Security Symposium (SEC '23)*. USENIX Association, USA, 5145–5162.
[50] Greg Kroah-Hartman. 2021. https://lore.kernel.org/linux-nfs/YH/fM/TsbmcZzwnX@kroah.com/.
[51] Heidi Laine. 2017. Afraid of scooping–case study on researcher strategies against fear of scooping in the context of open science. *Data Science Journal* 16 (2017), 29–29.
[52] Dana J Lawrence. 2007. The four principles of biomedical ethics: a foundation for current bioethical debate. *Journal of Chiropractic Humanities* 14 (2007), 34–40.
[53] Ruipeng Lei, Xiaomei Zhai, Wei Zhu, and Renzong Qiu. 2019. Reboot ethics governance in China. *Nature* (May 2019). http://www.nature.com/articles/d41586-019-01408-y
[54] Thomas A Limoncelli, Tom Limoncelli, Christina J Hogan, and Strata R Chalup. 2007. *The practice of system and network administration*. Pearson Education.
[55] Kevin Macnish and Jeroen van der Ham. 2020. Ethics in cybersecurity research and practice. *Technology in Society* 63 (2020), 101382.
[56] Kevin Macnish and Jeroen van der Ham. 2021. Ethical Approaches to Cybersecurity. *Oxford Handbook of Digital Ethics* (2021).
[57] John W. Meyer and Brian Rowan. 1977. Institutionalized Organizations: Formal Structure as Myth and Ceremony. *Amer. J. Sociology* 83, 2 (1977), 340–363.
[58] Ministry of Education, Sweden. 2003. Act (2003: 460) on ethical review of research involving humans (translated). https://www.riksdagen.se/sv/dokument-lagar/dokument/svensk-forfattningssamling/lag-2003460-om-etikprovning-av-forskning-som_sfs-2003-460.
[59] Arvind Narayanan and Bendert Zevenbergen. 2015. No encore for encore? ethical questions for web-based censorship measurement. *Ethical Questions for Web-Based Censorship Measurement (September 24, 2015)* (2015).
[60] Stuart G Nicholls, Tavis P Hayes, Jamie C Brehaut, Michael McDonald, Charles Weijer, Raphael Saginur, and Dean Fergusson. 2015. A scoping review of empirical research relating to quality and effectiveness of research ethics review. *PloS one* 10, 7 (2015), e0133639.
[61] Martin A Nowak and Karl Sigmund. 1992. Tit for tat in heterogeneous populations. *Nature* 355, 6357 (1992), 250–253.
[62] Sven Nyholm. 2023. *This is technology ethics: an introduction*. John Wiley & Sons, Inc, Hoboken.
[63] IEEE Board of Directors. 2014. IEEE CODE OF CONDUCT. *IEEE Policies* (2014). https://www.ieee.org/content/dam/ieee-org/ieee/web/org/about/ieee_code_of_conduct.pdf
[64] Office of Research Support and University of Texas Compliance. 2022. IRB submission forms. https://research.utexas.edu/ors/human-subjects/submitting-to-irb/irb-submission-forms/
[65] United States Government Accountability Office. 2023. *INSTITUTIONAL REVIEW BOARDS: Actions Needed to Improve Federal Oversight and Examine Effectiveness*. Technical Report. United States Government Accountability Office. https://www.gao.gov/assets/d23104721.pdf
[66] FDA Office of the Commissioner. 2023. https://www.fda.gov.
[67] Joseph P Simmons, Leif D Nelson, and Uri Simonsohn. 2021. Pre-registration: Why and how. *Journal of Consumer Psychology* 31, 1 (2021), 151–162.
[68] Sangeetha Paramasivan, Philippa Davies, Alison Richards, Julia Wade, Leila Rooshenas, Nicola Mills, Alba Realpe, Jeffrey Pradeep Raj, Supriya Subramani, Jonathan Ives, Richard Huxtable, Jane M. Blazeby, and Jenny L. Donovan. 2021. What Empirical Research Has Been Undertaken on the Ethics of Clinical Research in India? A Systematic Scoping Review and Narrative Synthesis. *BMJ Global Health* 6, 5 (May 2021), e004729. https://doi.org/10.1136/bmjgh-2020-004729
[69] Reethika Ramesh, Ram Sundara Raman, Apurva Virkud, Alexandra Dirksen, Armin Huremagic, David Fifield, Dirk Rodenburg, Rod Hynes, Doug Madory, and Roya Ensafi. 2023. Network responses to Russia's invasion of Ukraine in 2022: a cautionary tale for internet freedom. In *32nd USENIX Security Symposium (USENIX Security 23)*. 2581–2598.
[70] P. Resnick. 2014. *On Consensus and Humming in the IETF*. RFC 7282. IETF. http://tools.ietf.org/rfc/rfc7282.txt
[71] Thomas Reydon. 2013. *Wissenschaftsethik: Eine Einführung*. Ulmer.
[72] Thomas AC Reydon and Paul Hoyningen-Huene. 2011. Philosophie und ihr Verhältnis zu den Einzelwissenschaften. In *Warum noch Philosophie?* De Gruyter, 127–273.
[73] Thomas Ristenpart and Patrick Traynor. 2022. Call for papers of the 44th IEEE Symposium on Security and Privacy. https://web.archive.org/web/20240209062337/https://sp2023.ieee-security.org/cfpapers.html
[74] Gabrielle Samuel and Elizabeth Buchanan. 2020. Guest editorial: ethical issues in social media research.

[75] Sebastian Schrittwieser, Martin Mulazzani, and Edgar Weippl. 2013. Ethics in security research which lines should not be crossed?. In *2013 IEEE Security and Privacy Workshops*. IEEE, 1–4.

[76] signed by Swedisch Researchers. 2023. Öppet Brev Till Sveriges Riksdagsledamöter: Etikprövningslagen Kräver Omedelbar Revision. https://docs.google.com/document/d/1XAEAoqEp0NfhD6C1e5_wyQy553Faa866JAepYCYiSb4

[77] Konstantinos Solomos, John Kristoff, Chris Kanich, and Jason Polakis. 2021. Tales of favicons and caches: Persistent tracking in modern browsers. In *Network and Distributed System Security Symposium*.

[78] Wenlong Sun, Olfa Nasraoui, and Patrick Shafto. 2020. Evolution and impact of bias in human and machine learning algorithm interaction. *Plos one* 15, 8 (2020), e0235502.

[79] Sveriges Riksdågen. 2004. Lag (2003:460) om etikprövning av forskning som avser människor.

[80] Swedish Ethical Review Authority. 2023. Guide to the Ethical Review of Research on Humans. https://etikprovningsmyndigheten.se/for-forskare/utbildningsmaterial/. Online; accessed July 2024.

[81] Urmila M. Thatte and Padmaja A. Marathe. 2017. Ethics Committees in India: Past, Present and Future. *Perspectives in Clinical Research* 8, 1 (2017), 22–30. https://doi.org/10.4103/2229-3485.198549

[82] Petter Törnberg. 2023. Chatgpt-4 outperforms experts and crowd workers in annotating political Twitter messages with zero-shot learning. *arXiv preprint arXiv:2304.06588* (2023).

[83] Nancy Tuana. 2010. Leading with ethics, aiming for policy: new opportunities for philosophy of science. *Synthese* 177, 3 (Dec. 2010), 471–492. https://doi.org/10.1007/s11229-010-9793-4

[84] Jeroen Van Der Ham. 2017. Ethics and Internet measurements. In *2017 IEEE Security and Privacy Workshops (SPW)*. IEEE, 247–251.

[85] Jeroen van der Ham et al. 2015. Embedding ethics in system administration education. *The USENIX Journal of Education in System Administration* 1, 1 (2015), 1–9.

[86] Jessica Vitak, Nicholas Proferes, Katie Shilton, and Zahra Ashktorab. 2017. Ethics Regulation in Social Computing Research: Examining the Role of Institutional Review Boards. *Journal of Empirical Research on Human Research Ethics* 12, 5 (Dec. 2017), 372–382. https://doi.org/10.1177/1556264617725200

[87] Jessica Vitak, Katie Shilton, and Zahra Ashktorab. 2016. Beyond the Belmont Principles: Ethical Challenges, Practices, and Beliefs in the Online Data Research Community. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*.

[88] Hella von Unger and Dagmar Simon. 2016. *Ethikkommissionen in den Sozialwissenschaften-Historische Entwicklungen und Internationale Kontroversen*. Technical Report. RatSWD Working Paper.

[89] Qiushi Wu and Kangjie Lu. 2021. On the feasibility of stealthily introducing vulnerabilities in open-source software via hypocrite commits. *Proc. Oakland* (2021), 17.

[90] Qiushi Wu and Kangjie Lu. 2021. Withdrawal letter. https://www-users.cse.umn.edu/~kjlu/papers/withdrawal-letter.pdf

[91] Zhang Zhihao. 2022. New guideline will enhance oversight of research in frontier fields. http://en.qstheory.cn/2022-03/25/c_728125.htm.